Project Plan

Remote Exploration and Experimentation (REE) Project

March 1999

Agreements:

Edward C. Stone Director NASA Jet Propulsion Laboratory	Date
Eugene L. Tu HPCC Program Manager NASA Ames Research Center	Date
Robert D. Ferraro Project Manager NASA Jet Propulsion Laboratory	Date

Concurred by:

Barbara Wilson | Date
Director
Center for Space Microelectronics Technology
Jet Propulsion Laboratory

Michael J. Sander | Date
Director
Technology and Applications Programs
Jet Propulsion Laboratory

Joseph Bredekamp | Date
HPCC Executive Committee
NASA Office of Space Science

TABLE OF CONTENTS

Foreword

This document contains the Project Plan for the NASA Remote Exploration and Experimentation (REE) Project. This document is updated annually, and is the controlling document that defines the technical and management structure of the Project. The Project described in this document will accelerate the development of high-performance computing technologies to meet the needs of the spaceborne research community. It will also accelerate the distribution of these technologies to the American public. The technologies developed under this plan will help maintain U.S. technical and economic leadership in the international arena of high-performance computing. The time period covered by this plan is fiscal years 1999-2004.

The REE Project is a component of the NASA High Performance Computing and Communications (HPCC) Program, which in turn is part of the Federal program in Computing, Information and Communications (CIC). The primary goal of the Federal CIC effort is to extend U.S. technological leadership in high performance computing and computer communications. As this is accomplished, these technologies will be widely disseminated to accelerate the pace of innovation and improve national economic competitiveness, national security, education, health care, and the global environment. The NASA HPCC program is a critical element of the Federal CIC effort.

NASA's primary contribution to the Federal program is its leadership in the development of applications software and algorithms for massively parallel computing systems which will increase system performance to the sustained TeraFLOPS (10^{12} floating point operations per second) level for NASA applications. As HPCC technologies are developed, NASA will use them to solve its Grand Challenge research problems, whose solutions require significant increases in computational power and are critical to meeting national needs. NASA's Grand Challenges include improving the design and simulation of advanced aerospace vehicles, enabling people at remote locations to communicate more effectively and share information, increasing scientists' abilities to model the Earth's climate and forecast global environmental trends, and improving the capabilities of advanced spacecraft to explore the Earth and solar system. An additional component of the HPCC supports research and development of technology in education. Underlying and supporting all NASA program components is an element of basic research, which contributes, also, to providing new capabilities for the Next Generation Internet (NGI).

HPCC is a research program that pursues computing and communications technologies at various levels of maturity. It is structured to contribute to the broad Federal CIC effort, while addressing agency-specific computational problems that are beyond projected near-term computing capabilities. Computational problems in the areas of Earth science, space science, and aerospace are used as drivers of this research, providing the context and requirements for the work that is to be done. This work—and the HPCC Program—is organized into three Grand Challenge Computing Projects, a high performance communications project, and an education project.

- Computational Aerosciences (CAS)
- Earth and Space Sciences (ESS)
- Remote Exploration and Experimentation (REE)

- NASA Research and Education Network (NREN)
- Learning Technologies (LT)

These Projects, and their associated applications, were chosen for their potential and direct impact to NASA, their national importance, and the technical challenge they give the NASA HPCC Program. The document describing this program is the *Program Plan, High Performance Computing and Communications (HPCC)*.

1 Introduction

The Remote Exploration and Experimentation (REE) Project (UPN 626) is one of five Projects in the High Performance Computing and Communications (HPCC) Program. The Program is governed by the HPCC Program Commitment Agreement. Begun in 1992 as one of the original three Projects in the HPCC Program, it was deferred from 1993 – 1996 due to budget constraints. The Jet Propulsion Laboratory is the Lead Center for the REE Project. At this time, the Goddard Space Flight Center is supporting the REE Project by providing two science application teams.

1.1 History

"The REE element addresses critical needs to both the Offices of Space Science and Mission to Planet Earth. A new generation of on-board computers will enhance science return, reduce operations costs, and mitigate downlink limitations" [1]

Wesley T. Huntress, Jr.
Associate Administrator for Space Science

Charles F. Kennel
Associate Administrator for Mission to Planet Earth

It was with these prescient words in mind that the Workshop on Remote Exploration and Experimentation (REE) was convened in Pasadena, CA on August 21–23, 1995. The Workshop was followed by a Study Phase that took place from April–November, 1996. During this period, the REE Project consulted with US leaders in spaceborne avionics, High Performance Computing, commercial computing manufacturers, other government agencies, and NASA Space and Earth Scientists to devise a strategy and approach for meeting its Program Commitment Agreement (PCA) milestone in September 2003*: Demonstration of spaceborne applications on embedded high-performance computing testbed.* Key technical issues were examined, including: the current state of the art in spaceborne embedded computing systems, the trends in technology development for both spaceborne and commercial ground-based computing systems, and the projected computing requirements for several classes of NASA missions in the next millennium. Based on the results of the Study Phase, the REE Project developed a Vision and a set of Goals and Objectives which define the Project and its expected outcome. From these Goals and Objectives, a Level 1 Schedule of Milestones has been developed which will lead to demonstration of NASA spaceborne applications on a High Performance embedded computing system in space.

1.2 Vision, Goals

The commercial computing industry is two orders of magnitude larger than the entire space and defense electronics industry, and each year this disparity grows larger. The government no longer

---

[1] *Letter to: R/Director, High Performance Computing and Communications Office, July 6, 1994*

is a driving force in the state-of-the-art development of computing technology, and has little influence over its direction. At the same time, NASA and DOD requirements for space-capable computing technology are becoming more demanding, especially with regard to available power and cooling, performance, reliability, and cost. The REE Project seeks to leverage the considerable investment by the ground based computing industry to bring supercomputing technologies into space within the constraints imposed by that environment. The availability of onboard computing capability will enable a new way of doing science in space at significantly reduced overall cost. The vision of the REE Project, therefore, is:

To bring commercial supercomputing technology into space, in a form which meets the demanding environmental requirements, to enable a new class of science investigation and discovery.

Derived from this vision, REE has identified two principal goals. Specifically, the REE Project will:

Demonstrate a process for rapidly transferring commercial high-performance computing technology into ultra-low power, fault tolerant architectures for space.

Demonstrate that high-performance onboard processing capability enables a new class of science investigation and highly autonomous remote operation.

The legacy of the REE Project will not only be a new generation of scalable onboard supercomputing in space, but the validation of a process which will keep spaceborne computing capabilities on the same technology track as the commercial computing industry.

1.3 Overall Approach/Timeframe

Three Initiatives. Based on the results of the Study Phase, the REE Project developed a Technology and Applications Roadmap that leads to the attainment of the Project's goals and objectives and the accomplishment of the Project's PCA milestone in 2003. This roadmap is shown in Figure 1. It consists of three parallel interdependent initiatives, supporting the development of: (a) computing testbeds, (b) system software, and (c) applications. These initiatives work in concert with each other, merging in the end to achieve a spaceborne demonstration of scalable applications on a high-performance computing testbed.

Computing Testbeds Initiative. The purpose of the Computing Testbeds initiative is to explore and develop a process for translating commercial high performance scalable parallel computing architectures into low power spaceborne implementations. This architecture must rely, to the maximum extent practical, on commercial-off-the-shelf technologies and must minimize or eliminate the use of radiation-hardened components. The process must be consistent with the rapid (18 months or less) transfer of new earth-based technologies to NASA space missions. Translated architectures must satisfy a number of additional criteria, including no single point of failure and graceful performance degradation in the event of hardware failure.

The Computing Testbeds initiative will develop a series of hardware prototypes, leading to the demonstration of a capability of at least *300 MOPS*[2]*/watt.* This represents an increase of two

---

[2] *MOPS: Millions of Operations Per Second. These may be a mixture of 32 bit integer and floating point arithmetic or logical operations. Although MIPS (Millions of Instructions per Second) is a more traditional*

orders of magnitude over the power performance of the flight computer onboard the Mars Pathfinder spacecraft which landed on Mars in July, 1997. At the present time, a hardware testbed is being developed to demonstrate that significant power performance *(30 MOPS/watt)* can be achieved in a scalable embedded architecture using commercial technology. This testbed will also be the platform for conducting software implemented fault tolerance experiments and for developing the system software needed to achieve the reliability goals. The next step will be the design and fabrication of a hardware prototype which will match the mass and form factor of a future flight model and will demonstrate scalability *(50 nodes),* reliability *(0.99 over five years),* and a power performance of at least *300 MOPS/watt.*

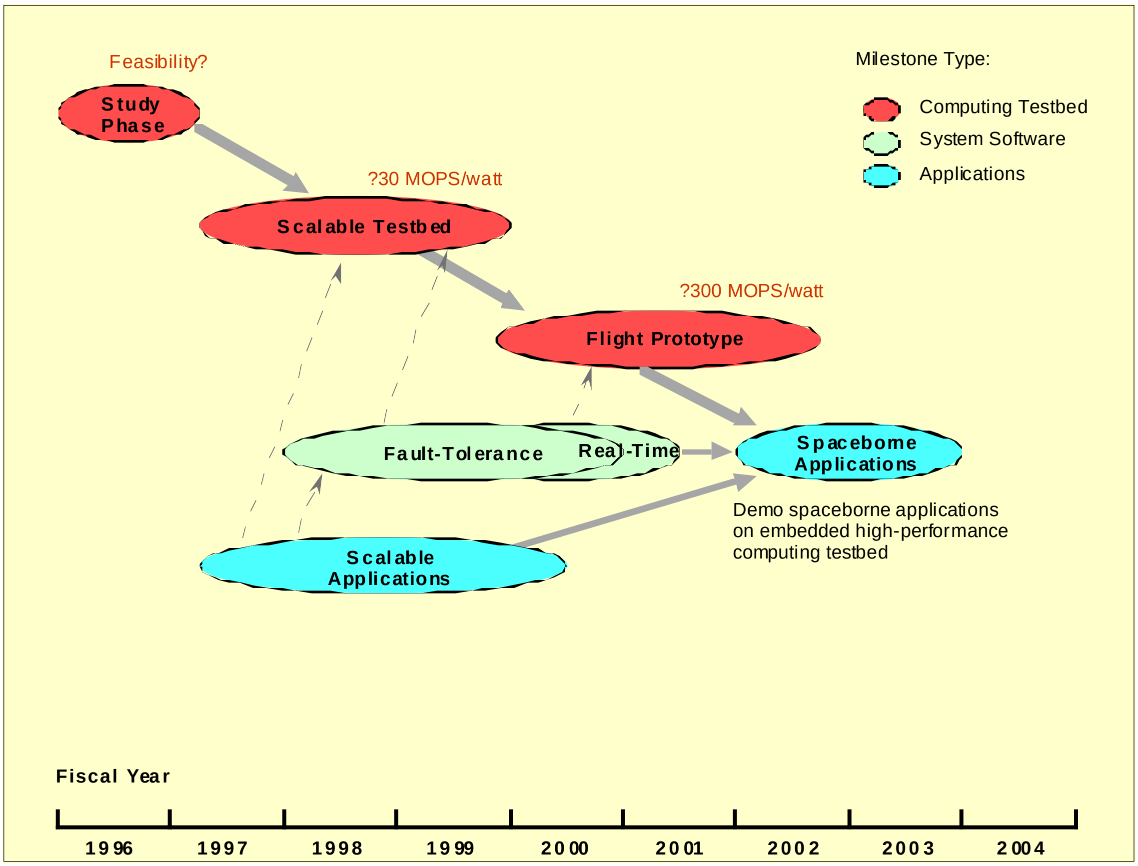

Figure 1. Technology and Applications Roadmap for the REE Project. This roadmap calls for the parallel development of hardware testbeds, systems software, and applications.

System Software Initiative. The purpose of the System Software initiative is to provide a set of services that will enable applications to take full advantage of the computing capacity of the hardware architecture, while providing an easy-to-use development environment and assuring reliable operation in space. By relying to the maximum extent practical on commercial software components, the system software layer will provide for the requisite performance capability and user interface. However, no commercially available parallel processing system offers a significant

---

measure of processor capability, it does not quantify the actual amount of work accomplished on processors which have complex instruction sets. In many cases, however, MOPS and MIPS will be interchangeable level of fault-tolerance without substantial task replication. Since the hardware architecture will be based on commercially available components, radiation-induced faults will be common and hardware component failure will be a possibility. Hence the system software must provide mechanisms for recovery from both permanent and transient faults. It will be a major challenge to the System Software initiative to develop a fault detection and recovery scheme that assures system reliability without compromising the performance capability available to the applications.

The System Software initiative will develop a middleware layer between a commercial operating system and the applications. This middleware layer will offer a suite of fault tolerance mechanisms from which the applications can make selections based on their reliability and efficiency requirements. The first version of the middleware layer will demonstrate reliability based on software implemented fault-tolerance *(0.99 over 5 years),* scalability *(50 nodes),* and portability for all REE applications. A later revision will add real-time capability as a feature.

Applications Initiative. The purpose of the Applications initiative is to demonstrate that the unique high-performance low-power computing capability developed by the Project enables new science investigation and discovery. Science Application teams will demonstrate that substantial onboard computational capability will be a crucial ingredient in science investigations of the future. They will ensure that architectures and system software produced by the Project meet the needs of the spaceborne applications community. They will stimulate the development and implementation of new computational techniques that will transform the REE platforms from computers into tools of scientific discovery, on a par with the sensors and data collection systems with which they are integrated.

The Applications teams will develop scalable science and autonomy application algorithms. Software will be developed and installed on the hardware testbed. This software will be used to test, evaluate, and validate candidate architectures and system software using the REE testbed. A demonstration of scalable applications on the hardware testbed will take place within months of its delivery. Subsequent generations of scalable applications for installation on the REE flight computer prototype will build on the experience gained in the hardware testbed. These applications will be demonstrated on the flight hardware prototype. The Project anticipates that there will be several flight opportunities available for this demonstration in the fiscal year *2002-2003* time frame

Raison d'être. These parallel tracks for the development of technology and applications are of equal importance. The significance of this point cannot be overemphasized. It is in the delivery of more science at lower cost that REE finds its ultimate *raison d'être.* This has motivated both an involvement in Space Science Enterprise and Earth Science Enterprise long-term planning and the creation of the Applications branch on the REE roadmap. It is through the involvement of users that the Project will introduce scalable spaceborne computing to the space science and autonomy communities and unearth the new mission concepts enabled by REE.

2 Objectives

From the Project Vision and Goals, REE has developed four specific Objectives:

1. Demonstrate power efficiencies of 300–1000 MOPS per watt in an architecture that can be scaled up to 100 watts, depending on mission needs.

2. Demonstrate new spaceborne applications on embedded high-performance computing testbeds which return analysis results to the earth in addition to raw data.
3. Develop fault-tolerant designs that will permit reliable operation for 10 years and more using commercially available or derived components.
4. Investigate ultra-low power onboard computer systems which will help open the entire Solar System to exploration without the need for nuclear technology.

These objectives address key issues in response to spaceborne computing requirements for the future. From the HPCC Program heritage of scalable multiprocessor systems, REE has derived its reliance on the commercial computing investment to provide components and architectures which have the capability to address NASA's onboard computing needs. The translation of HPCC technology to space, however, requires the Project to address issues of power, fault tolerance, and reliability which are different from the concerns of ground based computing. In particular, the limited available onboard power, the lack of ability to repair or replace failed components, and the need to compute in an environment which produces transient faults define the 1st, 3rd, and 4th objectives of the Project. The 2nd objective is prompted by a historical reticence on the part of the science community to do anything more than compress the data collected before transmitting it to the ground. The REE Project intends to demonstrate the usefulness of high performance embedded computing technology for enhancing the science returned in the presence of limited bandwidth to the ground and restrictive communications latencies to the spacecraft.

These objectives have driven the planning of the Project Milestones, which define the path to be taken towards the achievement of these objectives. Each milestone has a set of performance metrics which define the required capability in hardware, applications performance, software reliability, and overall system performance. These metrics are defined later in this document along with the Project Level 1 Milestones in the **SCHEDULES** section.

3 Customer Definition and Advocacy

A fundamental goal of the REE Project is to enable the return to Earth of dramatically new science results and insight from NASA spacecraft, using the unique high-performance low-power spaceborne computing capability developed by the Project. REE is a technology push project, designed to inject the HPCC Program scalable computing technology into NASA's spaceborne exploration activities. The customer base REE seeks to satisfy is future NASA science missions that face severe constraints on onboard power, cost, and communications bandwidth and latency to the ground. REE technology will not initially be targeted for used in routine spacecraft control (e.g., attitude control and thruster firing), although there is nothing inherent in our approach that precludes this. Such control tasks are not computationally bound and can be managed by state-of-the-art single-string radiation-hardened processors. REE technology will initially be used to provide high throughput (with high availability) for data-prolific science instruments. It may in addition be used for "spacecraft control" in the sense that for some applications, onboard computing will enable real-time redirection of the observing program based on the identification of a science target of opportunity. The Project is chartered to take the risk in introducing the latest commercial technology into space, solving the reliability and implementation problems, and transferring the technology to the mainstream of NASA's space missions. In order to adopt this new technology, the mission customers must be convinced that their reliability is not compromised, their capability is enhanced, and their budgets are not impacted.

To achieve this goal, REE engages high profile mission scientists to lead its applications teams. It is the science mission principle investigator who will ultimately define the required science return, which in turn sets the requirements for spacecraft capability. The Project seeks to maximize that return for a given cost by enabling new scientific investigations supported by capable onboard computing. These investigations will be defined by our primary customers, the science research community, in particular by space and earth science instrument Principal Investigators.

Five teams of science and autonomy investigators have been assembled by the REE Project to put forth specific proposals for novel applications to exploit the scalable hardware and system software. These teams are listed in Table 1. They will perform the following crucial functions:

1) They will develop revolutionary new mission concepts that utilize substantial onboard computational power as a crucial ingredient in scientific data collection, analysis, editing, and discovery.
2) They will ensure that architectures and system software produced under the Project match the scientific needs of the spaceborne applications community.
3) They will drive the implementation of new algorithms and computational techniques that transform the REE platforms from computing devices to tools of scientific discovery, on a par with the sensors and data collection devices with which they are integrated.
4) They will form the nucleus of an extended community of advocates for the utilization of spaceborne computing as a tool for remote exploration and experimentation in the planning and execution of NASA missions.

Table 1. REE Science Application Teams

Application	Principal Investigator	NASA Theme Addressed
Gamma Ray Large Area Space Telescope (GLAST)	Prof. Peter Michelson Stanford University	Structure and Evolution of the Universe (SEU)
Mars Rover Science	Dr. Steven Saunders Jet Propulsion Laboratory	Exploration of the Solar System (ESS)
Next Generation Space Telescope (NGST)	Dr. John Mather Goddard Space Flight Center	Structure and Evolution of the Universe (SEU)
Orbiting Thermal Imaging Spectrometer	Prof. Alan Gillespie University of Washington	Earth Science Enterprise
Solar Terrestrial Probe	Dr. Steven Curtis Goddard Space Flight Center	Sun-Earth Connection (SEC)

Throughout the life of the Project, the set of science application teams will evolve to continue to cover the mission areas which are of importance to NASA. It is the advocacy of these missions which is crucial to the Project's success.

4 Project Authority

The overall project authority for the REE Project is established by the HPCC Program, which is in turn established by the NASA Headquarters Program Management Council. The HPCC Program Commitment Agreement (PCA) represents the Agency-level agreement for the implementation of the HPCC Program and its Projects. Although the program is funded by three Enterprises and the NASA Education Division, the overall management of HPCC is formally within the Aero-Space Enterprise and is the responsibility of the HPCC Program Office at the NASA Ames Research Center.

The Jet Propulsion Laboratory is the designated lead center for the Remote Exploration and Experimentation Project. JPL has REE Project Management authority and responsibility. The NASA Goddard Space Flight Center (GSFC) currently supports the REE Project through its participation in the development of algorithms and software for the Next Generation Space Telescope and the Solar Terrestrial Probe science application teams. Other NASA centers may be called upon from time to time to support specific development activities in the Project as needed.

5 Management

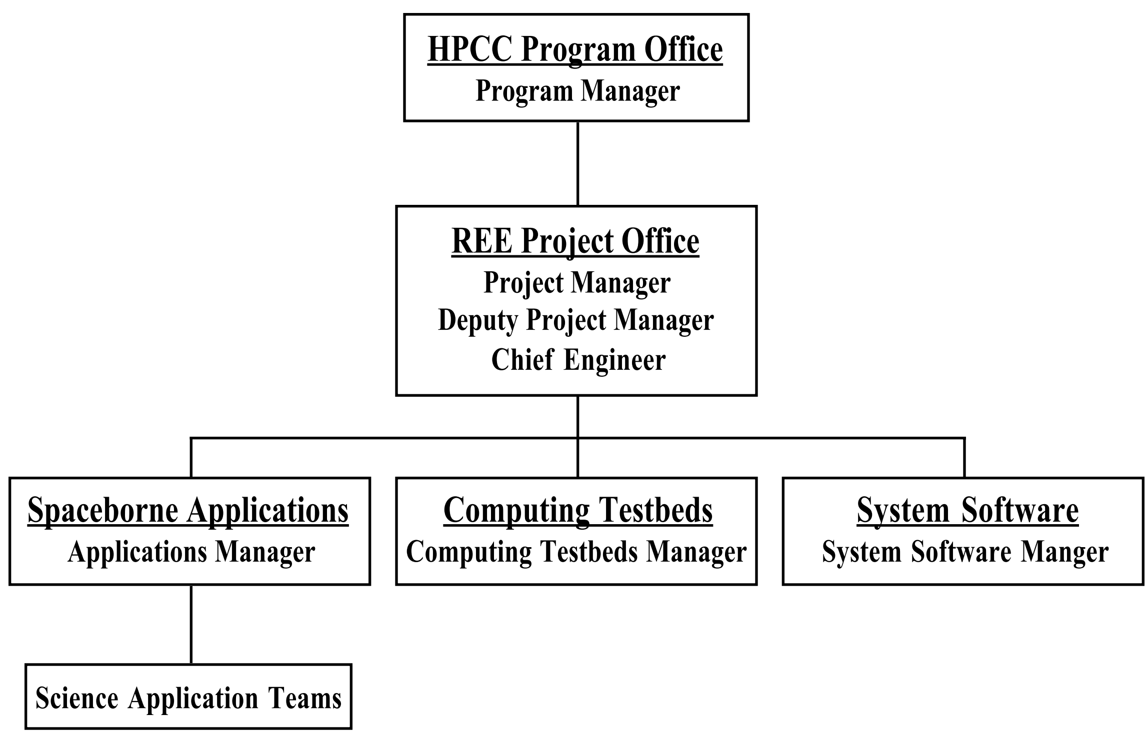

Figure 2. Management structure of the REE Project

5.1 Organization

The REE Project is managed by the REE Project Manager who reports to JPL Center Management and to the HPCC Program Manager. The REE Project Manager directs and controls the day-to-day activities necessary to accomplish Project goals and ensure customer satisfaction. Work performed at JPL for this project will be executed according to Policies and Procedures of the JPL Develop Needed Technology Domain. The REE Project Manager is assisted by a Deputy Project Manager, and a Chief Engineer. The major Work Breakdown Structure (WBS) elements are each lead by an element Manager who is responsible for the day to day activities within these areas. A suite of Applications have been identified and put under contract to the REE Project. Each Application has designated a Principal Investigator, who oversees day-to-day activities and reports to the REE Applications Manager. Figure 2 shows the management structure of the Remote Exploration and Experimentation Project.

5.2 Responsibilities

5.2.1 Project Manager

The overall management of the Remote Exploration and Experimentation Project is the responsibility of the REE Project Manager who is appointed from the Technology and Applications Programs Directorate at JPL. The specific responsibilities of the REE Project Manager are:

(a) Develop, update, and maintain the REE Project Plan, including the definition and negotiation of resource, schedule, and deliverable commitments, in cooperation with functional managers in the participating and sponsoring organizations.

(b) Direct and control the day-to-day activities necessary to accomplish the goals and objectives of the Project and to ensure customer satisfaction.

(c) Coordinate REE activities with those of the other HPCC Projects and participate through the HPCC Program in the Federal High Performance Computing and Communications Program.

(d) Coordinate REE activities with those of related programs in other government agencies, such as the Defense Advanced Research Projects Agency (DARPA).

(e) Appoint the Deputy Project Manager, the Chief Engineer, and element Managers for each of the three major REE WBS elements (Applications, Testbeds, and System Software), define and interpret element area responsibilities and statements of work, and direct and coordinate their efforts.

(f) Achieve all Project milestones.

5.2.2 Deputy Project Manager

The Deputy Project Manager (DPM) is appointed by the Project Manager and assists the Project Manager in the day-to-day operation of the project.

5.2.3 Chief Engineer

The Chief Engineer is appointed by the REE Project Manager. The specific responsibilities of the Chief Engineer are:

(a) Organize and coordinate the technical activities of the Project.

(b) Review and approve all REE technical documentation including the documentation tree, plans, procedures, schedules, specifications, requirements, level 3 milestones, and reports.

(c) Integrate and test all Project deliverables, with the support of the element Managers of the three major REE WBS elements.

(d) Work with the Project and WBS element Managers to define and maintain an integrated schedule of all Project activities, conduct fault and risk management studies, and define all test procedures and acceptance criteria for Project deliverables.

5.2.4 Applications Manager

The Applications Manager is appointed by the REE Project Manager. The specific responsibilities of the Applications Software Manager are:

(a) Define (with the concurrence of the Chief Engineer) the Level 2 and Level 3 milestones and schedules necessary to achieve the REE Level 1 Applications Milestones.

(b) Manage the budget for all Applications Software activities.

(c) Provide technical assistance to the Application Teams as needed

(d) Represent the Application Teams to the Project and to the system software activities.

(e) Provide the applications necessary to achieve all Project milestones.

5.2.5 Computing Testbeds Manager

The Computing Testbeds Manager is appointed by the REE Project Manager. The specific responsibilities of the Computing Testbeds Manager are:

(a) Define (with the concurrence of Chief Engineer) the Level 2 and Level 3 milestones and schedules necessary to achieve the REE Level 1 Computing Testbeds Milestones.

(b) Manage the budget for testbed acquisitions and operations.

(c) Provide the computing testbeds hardware necessary to achieve all Project milestones.

5.2.6 System Software Manager

The System Software Manager is appointed by the REE Project Manager. The specific responsibilities of the System Software Manager are:

(a) Define (with the concurrence of Chief Engineer) the Level 2 and Level 3 milestones and schedules necessary to achieve the REE Level 1 System Software Milestones.

(b) Manage the budget for all System Software activities.

(c) Oversee the design, implementation, and testing of software implemented fault tolerance layers.

(d) Provide the system software necessary to achieve all Project milestones.

5.2.7 Field Center Responsibilities

The Jet Propulsion Laboratory is the lead center for the REE Project. JPL will provide the technical lead and Project Management for REE. Discussions with other NASA Field centers may in the future lead to accountable roles for these centers to play in the Project.

5.2.8 Reporting Responsibilities

The REE Project Manager will submit status management and financial reports to the HPCC Program Manager as specific in the HPCC Program Plan. On an annual basis, the REE Project Manager will prepare an accomplishments summary suitable for inclusion in the HPCC Annual Report.

5.2.9 Coordination with Related Programs

The REE Project will coordinate its activities with those of related programs in other government agencies. In particular, REE is closely coordinating its activities with Air Force Research Laboratory (AFRL) Improved Space Computer Program (ISCP).

6 Technical Summary

The Technical Summary is divided into four major subsections. These are: Applications (section 6.1), Computing Testbeds (section 6.2), System Software (section 6.3), and Advanced Technology Opportunities (section 6.4). The relationships among the first three major Project activities and the strategy behind their structure is shown in figure 1. Advanced Technology Opportunities are high risk, high payoff investments which potentially cross-cut the first three activities and could result in REE significantly exceeding its Project goals, should they be successful. However, these investments are not on the Project's critical path.

The Project has been constructed in three distinct phases:

In the study phase, it was determined that the Project was feasible and quantitative goals were defined.

In the scalable testbed phase, a first testbed is being built for the purpose of proving out the basic principles as well as to perform initial development work and to explore the system trade space. During this phase, a hardware test bed is being fabricated in the currently available state of the art technology, system software is being developed for the testbed, various approaches to software implemented fault tolerance from both the academic and commercial sectors are being validated, and an initial set of representative applications are being ported to the testbed. In this experimental, proof of principle phase, an initial implementation of a REE computer is fabricated, its limits architectural features are explored, improvement in power performance is tested and validated, and new applications that can be enabled with an REE computer are validated and demonstrated.

In the flight prototype phase, the lessons learned from the scalable testbed phase are used to create a protoflight system which is form fit and function flight ready. In this phase, a flight ready system with state of the art hardware and software components in an optimized architectural configuration is fabricated. Additional science teams to expand the range of applications and broaden the new-science thrust of the Project will also be engaged. The final system will be tested, validated and qualified for flight.

The final result of these three phases will be a flight system demonstration and, potentially, a flight experiment in the '02-'04 time frame.

The following sections detail each of the major activities, i.e., Applications, Computing Testbeds, System Software, and Advanced Technology Opportunities, and their relationship to the overall Project strategy and organization.

6.1 Applications

A fundamental goal of REE is to enable the return to Earth of dramatically new scientific results and insight from NASA spacecraft, using the unique high-performance low-power spaceborne computing capability developed by the Project. To achieve this goal, new scientific directions will be defined by the scientific research community, especially by space and Earth science instrument Principal Investigators.

6.1.1 Science Strategy and Approach

Five teams of science and autonomy investigators have been assembled by the REE Project to put forth specific proposals for novel science applications to exploit the scalable hardware and system software. These teams will perform the following crucial functions:

1) Develop revolutionary new mission concepts that utilize substantial onboard computational power as a crucial ingredient in scientific data collection, analysis, editing, and discovery.
2) Ensure that architectures and system software produced under the Project match the scientific needs of the spaceborne applications community.
3) Drive the implementation of new algorithms and computational techniques that transform the REE platforms from computing devices to tools of scientific discovery, on a par with the sensors and data collection devices with which they are integrated.

The REE Applications will highlight entirely new ideas. There are two fundamental reasons for this. First, they will have available the unique resources supplied by REE: at least two orders of magnitude more computational power than has previously been available in space. Second, these resources may be deployed on miniature spacecraft orders of magnitude smaller than those currently in existence, with severely limited electrical power for data transmission to earth.

These two observations imply that REE-enabled space science and Earth observing missions of the future will look very different from any that have taken place to date. REE will focus heavily upon missions featuring autonomous spacecraft, either entirely self-commanded or able to operate with minimal supervision from the ground based upon the transmission of limited instructions and data. For example, an REE computer would enable vigilant spacecraft, or spacecraft fleets, that will be able to monitor planetary, Earth, solar and stellar targets continuously for weeks, months, even years at a time. Applications implemented on the REE computer will be able to flag hazardous or scientifically interesting events as they occur, allowing the spacecraft to respond autonomously, either to maintain its own health in the face of hazards, or to image especially interesting behavior at higher resolution so that the most scientifically important results can be returned to Earth.

Spacecraft autonomy is already a vigorous focus of future spacecraft planning at NASA. It is a major goal, for example, in JPL's New Millennium Program. The unique ingredient that will be provided by REE is the ability to pursue science-driven autonomy, which is currently considered only in research programs such as the Office of Space Science Autonomy and Operations Technology Program.

Note that the onboard computing capability provided by REE is absolutely crucial to the achievement of scientific goals in situations in which a rapid adaptive response to unexpected events is needed. For example, it may be required to capitalize on important transient activity in an imaged target. This point is often greatly underappreciated. It is typically assumed, for example, that the decision to rely upon substantial onboard computing to process scientific data depends solely upon the telemetry bandwidth available to transmit this data to Earth. However, in regimes such as deep space, there is often insufficient time to return data to Earth and to await further instructions during an interesting unexpected transient occurrence, *even if sufficient telemetry bandwidth is available.* In other cases, future competition for resources such as JPL's Deep Space Network (DSN) will severely limit the amount of downlink available to individual missions, even for spacecraft operating in relatively power-rich environments such as an orbit of Mercury.

Accomplishment of science-driven autonomy goals will require a suite of new algorithms and applications software to be developed as part of REE, to ensure that hardware capabilities of the REE computers are exploited to their fullest. These include automated onboard data analysis of remote sensing imagery, autonomous navigation and control software, planning and scheduling of resources, data compression and editing, and the construction of onboard catalogues and models as scientific reference points in the knowledge discovery process. These activities must be defined and prioritized by the science community as part of their involvement with REE.

Each Application team will have the following Project responsibilities:

1) Identify important new scientific directions which may be enabled by REE.

2) Analyze the computational requirements, especially with respect to CPU, RAM, I/O, sophistication of programming model, importance of fault tolerance, and operating system needs. These analyses, to be completed in the first six months of science team contracts, will be used to evaluate the testbed architecture.

3) Develop algorithms and prototype applications on ground testbeds to demonstrate feasibility. Initially, access will be provided to traditional HPC platforms such as

the Cray T3E, SGI Origin, and HP Exemplar machines, supporting parallel API's such as MPI, to serve as emulators. These will be augmented by ground testbeds provided at a later date, as described in the Computing Testbeds section in this plan.

4) Explore new approaches to science on board in a limited downlink bandwidth environment.

The applications teams will become an integral part of the Project, and will be ultimately responsible for the Project's applications demonstration milestones as outlined in Appendix B.

6.1.2 Selection of Application Teams

Five application teams were selected in fiscal years 1997/1998 to participate in the REE Project. They are: (i) Gamma Ray Large Area Space Telescope (GLAST), (ii) Next Generation Space Telescope (NGST), (iii) Mars Rover Science, (iv) Orbiting Thermal Imaging Spectrometer (OTIS), and (v) Solar Terrestrial Probe Multiplatform Missions. These teams are lead by NASA scientists. They are developing algorithms and software for applications that emphasize *in-situ* analysis of science instrument data and remote operation of highly autonomous systems. These Applications were chosen on the basis of their potential for benefiting from the hundred-fold increase in onboard computing power that REE promises. The operation of their instruments is constrained because of the combination of their high data rates or limitations in spacecraft downlink bandwidth, or both. In some cases, operation is constrained by latency.

These five Applications, their science objectives, and the attributes that drive their science requirement, are described in detail in Appendix C.

6.2 Computing Testbeds

The purpose of the Computing Testbeds initiative is to transition a commercial scalable high performance computing architectures into forms which are appropriate for a spaceborne computer. This architecture must rely, to the maximum extent practical, on commercial-off-the-shelf technologies and must minimize or eliminate the use of radiation-hardened components. The approach must be consistent with the rapid (18 months or less) transfer of new earth-based technologies to NASA space missions. The architecture must satisfy a number of additional criteria, including no single point of failure and degrading performance gracefully in the event of hardware failure.

The Computing Testbeds (CT) initiative consists of three Level One milestones. The first of these (CT 5) called for a Study Phase. This was successfully executed and completed *(June - Nov. 1996),* establishing the feasibility of the Project's goals and objectives. The second milestone (CT 8) calls for the development of a hardware testbed that will demonstrate scalability *(50 nodes)* and power performance *(30 MOPS/watt).* This testbed completed its design phase in September 1998. It will be fabricated and delivered in November *1999.* The third Level One milestone of the CT initiative (CT 10) calls for the development of a hardware prototype. The prototype will match the mass and form factor of a future flight model and will demonstrate scalability *(50 nodes),* reliability *(0.99 over five years),* and a power performance of at least *300 MOPS/watt.* This represents an increase of two orders of magnitude over the state-of-the-art.[3] The hardware prototype will be delivered to JPL in *June 2002*.

In the following subsections, the First Generation Testbed (CT 8), the Flight Prototype (CT 10), and the development of ultra-low-power technologies are discussed in detail.

6.2.1 First Generation Scalable Embedded Computing Testbed (CT 8)

Beginning in fiscal year 1997, the REE Project formed a collaborative relationship with industry to develop a first-generation scalable, high-performance, low-power computing testbed. This testbed will be used to demonstrate scalability *(50 nodes)* and system-level power performance (at least *30 MOPS per watt)*. It will be used to test, refine, and validate scalable architectures and system approaches to fault tolerance prior to investing in the development of a flight prototype. This testbed is to be delivered to JPL in November 1999.

There are five key issues to be addressed and resolved during the testbed phase. These include:

1) *The development of scalable high-performance low-power architectures.* The performance of single-CPU flight computers has shown steadily increasing capability from Voyager *(7 KIPS, 0.1 KIPS/watt)* to Mars Pathfinder *(20 MIPS, 2.2 MIPS/watt).* The REE Project will develop and evaluate a scalable low-power high-performance multiprocessor *architecture* based largely on terrestrial computing experience, to improve onboard performance by *100x* or more over the next five years.

2) *The development of ultra-low-power components.* The REE Project will accelerate the development of efficient onboard computing by investing in ultra-low-power *component* technologies, such as processor-in-memory and reconfigurable computing. (See Section 6.2.3.)

3) *The development of COTS-based system-level fault- and radiation-tolerance.* The REE Project will support the development of system-level software and hardware fault-tolerance features that will allow the system to operate dependably in the high radiation environment of space, while making extensive use of COTS-based parts. The use of COTS-based technologies is necessary to assure user access to the latest commercial technologies, maximize performance, and reduce cost.

4. *The development of autonomous graceful degradation.* The REE Project will support the development of autonomous features that will allow the system to degrade gracefully and assure reliable operation over mission lifetimes of ten years and more.

5. *The development of a real-time capability in multiprocessor architectures.* The REE Project will support the development of an operating system and software development environment for scalable multiprocessor architectures that supports real-time operation.

---

[3] *Mars Pathfinder, July, 1997*

A solicitation was issued at the end of fiscal year 1997 inviting proposals from teams led by industry, or possibly academia, to develop a testbed platform to investigate scalable low-power high-performance architectures, based largely on COTS technologies. Following evaluation of the proposals received, contracts were awarded to two teams, led by: Sanders, a Lockheed-Martin Company, of Nashua, NH and SEAKR Engineering, Inc. of Englewood, CO. (Table 3.) These vendors were selected for a six-month Design Phase, at the end of which one would be selected to fabricate and deliver the hardware testbed. The testbed design is required to contain 20 nodes, at least four of which are fully functional hardware nodes, capable of demonstrating the power performance requirement. Although this testbed will contain only 20 nodes, it can be used to investigate scalability to larger configurations by a combination of experiment and analysis. The testbed will interface with the existing JPL Flight Systems Testbed, permitting the validation of the computer interfaces with typical spacecraft busses and other spacecraft subsystems, such as sensors, telecommunications, and controllers. Applications developed by the Applications Teams will be ported to the testbed to evaluate the total system performance for a variety of spaceborne computing scenarios.

Table 3. Participants in the REE Testbed Design Phase (3/98-9/98).

Lead Organization	Collaborating Organizations
Sanders, a Lockheed-Martin Company	California Institute of Technology Lockheed Martin Federal Systems University of Illinois University of Southern California
SEAKR Engineering, Inc.	Lockheed Martin Control Systems Lockheed Martin Tactical Defense Systems Motorola Corp. SGI/Cray Research, Inc.

At the end of the Design Phase, Design Reviews were held at the home facilities of the two lead organizations. Based on an evaluation of these proposals by a source evaluation team, Sanders was selected to fabricate and deliver the First Generation Testbed. Sanders was placed on contract with JPL in early November 1998. The First Generation Testbed will be delivered in November 1999.

6.2.2 Flight Prototype Embedded Scalable Computer (CT 10)

In fiscal year 2000, the REE Project will begin development of a prototype flight computer. The architecture of this platform will be based on the refinements developed through the experience with the first generation testbed. This prototype will match the mass and form factor of a future flight model and will demonstrate scalability *(50 nodes),* reliability *(0.99 over five years),* and a power performance of at least *300 MOPS/watt.* As with the testbed, the prototype will be developed in partnership with industry. Application software developed by the applications teams will be installed on this prototype. The prototype will be used to demonstrate a low power, scalable architecture capability using the latest generation COTS and low power component technologies with a systems level approach to fault tolerance and real-time capability. The prototype, application software, and associated instrument, will be deployed in space as a science and technology demonstration. The criterion for a successful demonstration will be the performance of science research or autonomous operations explicitly enabled by REE technology. The Flight Prototype Computer will be delivered in June 2002.

6.3 System Software

The primary goal of the REE system software effort is to provide a set of services which enables applications to take full advantage of the computing capacity of the REE architecture while providing an easy-to-use programming and development environment. In addition, the system software must provide for fault detection and fault recovery so that applications can operate in the presence of faults. To the maximum extent possible, a commercial scalable multiprocessor operating system will be baselined and the added functionality will be layered on top of it.

The System Software (SS) initiative consists of two Level One milestones. The first of these calls for the development of a software-implemented fault-tolerance (SIFT) capability that will support system reliability *(0.99 over five years)*. In addition, the system software must support scalability *(ln or better/50% of ideal)* and applications portability. This capability will be demonstrated on the hardware testbed by *September 2000.* The second SS Level One milestone adds real time capability *(20 ms performance latencies)* as a requirement. This added capability must be demonstrated by *March 2001.*

6.3.1 User Access

A fundamental requirement of the REE-based system is that it be easy to use and support the needs of the user community. Ideally, the user will develop, validate, and update application software on his or her laboratory workstation. Updated software will then be installed on the REE platform and operated with the same user interface as on the workstation. To facilitate this goal, the system software will utilize tools, interfaces, and programming languages that are based on standards and commercial products that are familiar to the user community.

System Services. The REE hardware system will consist of a set of computing nodes and memory interconnected by a network fabric. Based on input from the REE applications teams, it appears that a relatively simple programming model (based on explicit message passing) will be sufficient for REE applications. In addition, the system software must provide task management to enable task creation, deletion, context switching, and scheduling. Also, the operating system on each node will provide memory protection facilities which will "fence off" applications from the operating system and each other, allowing multiple applications to run on any given node. The system software must also provide access to mass storage with an appropriate I/O model. The size of the prototype REE applications is relatively small, so it is not anticipated that paging or swapping functions will be required. A system software layer built from commercial software components should be able to provide these functions. In addition, the REE Project may investigate parallel languages and alternate programming models as candidate technologies for testbed validation.

Development Environment. The REE application teams have requested a complete set of applications libraries, providing functionality in linear algebra, signal and image processing, and statistics. Robust parallel versions of linear algebra routines have already been demonstrated which can detect errors in computation with virtually no increase in computational overhead. In addition, the applications teams have requested a set of parallel programming tools (e.g., performance monitors and debuggers) to assist in their code development. Commercially available parallel programming tools will be provided as part of the testbed delivery. In addition, the REE Project will coordinate with the CAS and ESS Projects to determine the applicability of the tools developed under these Projects to the REE environment.

6.3.2 Fault Tolerance and Real Time Operation

The REE architecture will be based on commercially-available, non-radiation-hardened components. In the space environments in which REE missions will operate, radiation-induced faults will be frequent, and the REE system must provide a mechanism for recovery from these faults. It is imperative that this recovery mechanism allow the applications to make progress in the presence of faults while still making efficient use of the available computing power.

In this context, it is important to distinguish between distributed applications and parallel applications. Both sets of applications are spread across multiple processors. Distributed applications are characterized by several cooperating tasks with multiple tasks per processor. These tasks are loosely coupled in the sense that they communicate infrequently, and the communication protocols that they employ can consist of several protocol layers without impacting the overall performance of the system. Parallel applications, by contrast, are characterized by several cooperating tasks in which a single task may be spread across multiple processors. The processors communicate almost continuously, and the communication protocols must be extremely efficient in order not to impact the overall performance of the system. Indeed, distributed systems may copy messages several times as messages are passed from one protocol layer to another, while parallel systems go to great lengths to avoid copying messages even once.

The REE Project is focused primarily on applications that are parallel. To date, most of the work in fault-tolerant multiprocessors has focused on distributed applications where fault tolerance can be implemented via relatively expensive mechanisms such as message duplication and task replication and voting, and relatively little attention has been paid to parallel applications. (Indeed, no commercially available parallel processing system offers any significant level of fault tolerance.) The challenge for the REE Project is to develop a system which provides fault tolerance with as little overhead as possible based on the reliability requirements of the application.

Real time operation is defined as the ability to respond to an external event, such as an externally generated signal to the system, and take appropriate action within a specified period of time. The time period must take into account the complexity of the response, but a guarantee of action within a defined period is what distinguishes real time systems from non-real time systems. The REE Project has set as its target a latency of no more than 20 milliseconds between the time an externally generated signal is input to the systems and the time at which the first instruction of the signal handler is executed. This latency is to be accomplished even in the presence of faults.

The REE Project will approach the development of software-implemented fault-tolerance and real-time capability primarily through partnering with industry. We will look to the private sector for cooperative development of a SIFT middleware layer to provide reliable operation on high-performance parallel hardware. In addition, the Project will leverage developments in a DARPA-sponsored program for the development of Distributed Reliable Computing. It is essential that the hardware and software be developed concurrently so that meaningful tradeoffs can be made during the design of both, resulting in the optimum system design.

6.4 Advanced Technology Opportunities

A fourth major effort targets the development of ultra-low-power processing and memory component prototypes. Several promising technologies exist which could have a high payoff in providing ultra-low-power systems. These technologies are all high-risk, because of their immaturity. The potential benefits, however, are quite large. Thus, REE will make a modest investment in an effort to develop an ultra-low-power computer using one of several high-risk/high-payoff technologies.

All of the low-power technologies under consideration owe their great promise to the same enabling technology trend: the ever shrinking size of solid state device features (e.g., transistors, conducting paths). This trend is leading to lower operating voltages (and hence lower power) and to a greater density of gates that can be placed on silicon. As more gates can be placed on a single chip, so also can more functionality.

Ultimately, this trend will permit fully functional general purpose computers (including multiple CPUs, RAM, an interconnect structure, and off-chip drivers) to be placed on a single chip. This approach, commonly referred to as processor-in-memory (PIM), has enormous advantages. First, the considerable power normally invested in moving data between chips (over 50% of the total in conventional architectures) is eliminated. Second, problems arising from memory bandwidth and latency, which invariably limit performance in conventional architectures, are dramatically reduced. Secondary benefits follow. For example, architectures may be simplified as the need for a complex cache structure is reduced. Possibly, caches will be eliminated altogether.

This trend will also permit the placement of high-functionality special-purpose computers on a single chip. The advantages of special purpose processors have been known for decades, with application-specific integrated circuits (ASICs) out-performing comparable-sized general-purpose processors by an order of magnitude or more. But ASICs come with an enormous disadvantage: They have no flexibility. Once fabricated and launched, an ASIC cannot be changed. About a decade ago, field programmable gate arrays (FPGAs) were introduced, which added general-purpose flexibility to ASIC performance. But still the low density of gates limited their functionality. As feature size continues to shrink, FPGA technology may be expanded, leading to a new generation of reconfigurable ASICs, gate arrays with sufficient capacity for use in the general purpose arena. The customization available through FPGAs may ultimately prove to provide the best overall power efficiencies for given level of computing capability. But substantial investment is required in the development of tools to program FPGAs using high level languages, so that digital logic designers can be eliminated from the programming and testing loop.

In fiscal years 1997 and 1998, the REE Project made an initial investment in the development of Processor-in-Memory (PIM) technology. Contracts were issued to Prof. Peter Kogge of Notre Dame to work with Lockheed Martin Federal Systems of Manassas, VA and insyte Corp. of Tampa, FL to develop and deliver a hardware PIM prototype to JPL. However, because of budget reductions in FY99, this effort was suspended. The REE Project plans to resume exploration of this promising technology area in fiscal year 2000.

Additional technology development opportunities may arise during the life of the Project. These may be software technologies as well as hardware technologies. The REE Project will periodically assess the return on its advanced technology investments and adjust its strategy as opportunities arise.

6.5 Summary

As seen from the above technical summary, the REE Project results in: a methodology for transitioning COTS components to space, a series of REE enabled science missions, and a first instantiation of an REE computer system ready for flight insertion. In addition it will have resulted in two generations of REE computers (Scalable Testbed and Flight Prototype) and several years of experience in working with both commercial vendors and mission scientists. The Project thus delivers a wealth of experience and proven capabilities by the time of first flight. It should also be noted that the REE Project should not end at the point of first flight insertion, but rather that a third generation and a fourth should follow in rapid succession to fully leverage the methodologies developed on the Project by continually keeping up with COTS state of the art computer systems. With each succeeding generation, system power performance will increase, while system development and fielding costs are reduced and newer, more powerful science missions enabled.

7 Schedules

The REE Project has defined a series of Level One milestones leading to the *demonstration of spaceborne applications on embedded scalable computing testbed,* as called for in the NASA/HPCC PCA. These milestones are listed in Table 4. The completion dates for several of these milestones have been changed from that of the previous version of this plan (May 1997). These changes were made in response to funding reductions imposed in fiscal years 1999 and 2000. A flight opportunity for a spaceborne demonstration of the final REE milestone GC 8 has not yet been identified. The Project anticipates that there will be several opportunities available in the fiscal year 2002 - 2003 time frame for this demonstration. Launch costs associated with a flight demonstration are not in the current budget requirements and are expected to be borne by another project (as yet unidentified). Metrics to be applied for each milestone are defined in Appendix B. The designation for each milestone is defined by its sequence within the entire set of Level One Milestones for the HPCC Program.

The REE Project Manager approves and maintains an integrated Level 2 Schedule developed by the REE Chief Engineer in consultation with the WBS element Managers. HPCC Program Manager concurrence is obtained for this schedule, and any changes to it during the course of Project execution. WBS element Managers develop and maintain lower level schedules as needed. The Chief Engineer approves these schedules.

Table 4. Level One milestones for the REE Project. Milestone types include Grand Challenge (GC), Computing Testbed (CT), and System Software (SS).

	Milestone	Date Due	Metrics	Performance Goals
CT 5	Complete studies of technology projections for embedded scalable high performance computing architectures in space	12/96 Comp 12/96	number of studies	5 Studies
CT 8	Install 1st generation scalable embedded computing testbed operating at 30-200 MOPS/watt	12/99	scalability	> 50 nodes
			performance/ watt	30 MOPS/watt in a flight configuration
GC 6	Demonstrate scalable spaceborne applications on 1st generation embedded computing testbed	03/00	number of applications	3 applications
			scalability/ speedup	*ln* or better/ 50% of ideal
			performance	50% of peak for architecture
SS 5	Demonstrate software implemented fault tolerance on 1st generation embedded computing testbed	9/00	scalability	*ln* or better/ 50% of ideal
			applications portability	all current REE challenge applications
			reliability	.99 over 5 years
SS 6	Demonstrate real time capability with software implemented fault tolerance for embedded scalable computers	3/01	scalability	*ln* or better/ 50% of ideal
			applications portability	all current REE applications
			reliability	.99 over 5 years
			real time performance latencies	20 milliseconds

CT 10	Demonstrate flight prototype embedded scalable computer operating at 300 - 1000 MOPS / watt	6/02	scalability	> 50 nodes
			performance/ watt	300 MOPS/watt in a flight configuration
			system reliability	.99 over 5 years
GC 8	Demonstrate spaceborne applications on embedded high-performance computing testbed	9/03	number of applications	3 applications
			scalability/ speedup	*ln* or better/ 50% of ideal
			performance/ watt	50% of peak for architecture
			reliability	.99 over 5 years

8 Resources

This section has been removed from the on-line version of this document.

9 Controls

9.1 Project Plan Changes

The process for controlling changes to the REE Project and the subordinate WBS elements is hierarchical and described in this section.

The Program Commitment Agreement (PCA) is the overall controlling document for the REE Project and the HPCC Program. It is a contract between the NASA Administrator and the Associate Administrator of Aero-Space Technology (AT), defining the high level requirements and commitments for the HPCC Program and the REE Project. Any changes to the REE Project that would affect the PCA, such as PCA milestones, must be agreed to by the HPCC Program Manager and approved by the AT Associate Administrator and by the NASA Administrator.

Any changes to the REE Project that do not affect the PCA but that do affect the HPCC Program Plan must be approved by the HPCC Program Manager and by the Office of the Associate Administrator for Aero-Space Technology.

For changes to the REE Project within the objectives, technical scope, schedule and budgets established in the approved Project Plan, the REE Project Manager has the authority to approve such changes.

A formal process is used for managing Project changes: requesting, acquiring the required level of approval, and tracking and documenting the changes. The REE Project Manager maintains the Project change log.

9.2 Computing Testbeds

All participants of the REE Project must comply with the NASA policy on access to software, data, and testbed facilities. Access to the REE testbeds will be open to U.S. citizens and U.S. permanent resident aliens. Access to the REE testbeds by foreign nationals requires advanced approval regardless of whether the foreign national is approved for physical access to JPL. The JPL Legislative and International Affairs Office is responsible for the initial foreign national approval process.

9.3 Sensitive Technology

The Government, and Caltech, shall have unlimited rights to technical data and computer software produced in the performance of contracts issued by JPL under the NASA REE Project. Unlimited rights, as used above, means the right to use, disclose, reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly in any manner and for any purpose, and to have or permit others to do so. These unlimited rights extend to the use of technical data contained in proposals upon which such contracts are based. Technical data and computer software developed at private expense, including minor modifications thereof, remain the property of the developing entity and are protected from unauthorized disclosure and use. Government rights and the rights of Caltech (the California Institute of Technology) are defined by the JPL Prime Contract with NASA, which governs all activities undertaken by JPL.

All information released by JPL outside of JPL will be done in accordance with JPL Policy: *Releasing Information Outside of JPL.* The release to a foreign national of technical information that resides at or is controlled by JPL requires advanced approval through the JPL Legislative and International Affairs Office as described in Section 9.2 above. The REE Project will adopt conventional security techniques of isolating critical technology from "open" exchange systems until acceptable multi-level security techniques and policies have been developed.

Negotiated License Agreements are used to restrict access to privately developed technology performed under the auspices of the REE Project. These agreements provide NASA with limited rights to use proprietary data or designs in NASA in-house or cooperative research projects. These agreements specify limits on the distribution and use of the proprietary data by NASA and NASA-licensed entities.

Some sensitive information developed solely within the REE Project may be subject to protection under the Export Administration Regulations (EAR) or the International Traffic in Arms Regulations (ITAR), which are export controls established by law. The participants in the REE Project will follow applicable export control laws. These regulations establish lists or categories of technical data and/or products that may not be exported without an approved export license. (Note that the definition of "exported" includes "disclosed" and "discussed" as well as published.)

Technical data and computer software produced for REE by another NASA Center is governed by each Center's Policies and Procedures for the control of Sensitive Technology. Work performed at other NASA Centers shall comply with that Centers Policies and Procedures, and applicable Federal Law.

10 Implementation Approach

The development of the Project hardware deliverables will be done largely out-of-house. RFPs were issued for the execution of a Study Phase and for the fabrication and delivery of a First Generation Testbed. An RFP will be issued for the eventual development of a flight prototype. The development of Project system software deliverables will be done at JPL in partnership with industry, academia, and other government agencies. External contracts will be submitted to competitive bidding to the maximum extent practical. The development of algorithms and software for applications will be lead by NASA scientists or mission managers, with the work performed largely at their home institutions.

10.1 REE WBS

A Work Breakdown Structure (WBS) has been developed to reflect the major activities being undertaken over the life of the Project. This WBS is organized around the principle technical activities of the Project, as detailed in **Technical Summary** section, and the cross-cutting functions of Project Management and System Engineering. The activities under Advanced Technology Investigations are fluid and additional activities may be added during the life of the Project, at the discretion of the Project Manager.

1.0 Spaceborne Applications

- 1.1 Science Applications
- 1.2 Applications Technical Support
- 1.3 Application-Based Fault Tolerance

2.0 Embedded Computing Hardware Research and Development

- 2.1 Testbeds
- 2.2 Early Prototype

3.0 System Software Research and Development

- 3.1 Software Implemented Fault Tolerance Architecture
- 3.2 Software Implemented Fault Tolerance Development
- 3.3 Prototype System Software

4.0 System Engineering

- 4.1 Studies
- 4.2 Modeling
- 4.3 System Design
- 4.4 Validation and Test

5.0 Advanced Technology Investigations

- 5.1 Processor In Memory (PIM)

6.0 Management

6.1 Project Management

10.2 Project Descope Process

Should descoping of the REE Project or rescoping of any of its constituent WBS elements be required, whether due to resource reductions in the REE Project or the need to rebalance the resources within the Project, the following descope process will be followed.

1. The REE Project Manager, in consultation with the Chief Engineer and element Managers, will develop a list of current Project activities on the critical path for Level One Milestones. Each element Manager will define the minimum level of activity required to adhere to schedule.
2. The Project Manager will rebalance the available resources to maintain schedule at the expense of increased risk of failure to achieve Level One milestones on time. Risk to testbed milestones will be increased first, application milestones second, and system software milestones last.
3. If schedule cannot be maintained with the available resources, the Project Manager will attempt to reschedule Level One Milestones to conform to the expected resources profile, and request concurrence from the HPCC Program Manager.
4. If rescheduling Level One Milestones is not possible under the expected resources profile, the Project Manager will propose a new set of Level One Milestones which correspond to reduction in demonstrated system capability at the end of the Project, and request concurrence from the HPCC Program Manager. Hardware performance will be targeted first, followed by real time SIFT capability.
5. In the event that the available resources no longer support the development of SIFT capable of handling the fault rates in low Earth orbit or deep space, the Project Manager shall recommend to Program Management that the REE Project be canceled.

11 Acquisition Summary

Free and open competitive procurements will be used to the maximum extent possible. The primary procurement vehicle that is expected to be put to use in the REE Project is the Request for Proposals (RFP). At JPL, this vehicle results in contracts. Interagency agreements for joint R&D endeavors may also be used as the occasion arises.

12 Project Dependencies

The planned spaceborne demonstration of an embedded HPC system is dependent on the identification of a flight opportunity in which the launch and operations costs are borne by some other project. These costs are not budgeted within REE. Failure to identify such an opportunity will require REE to meet its final milestone through a ground-based demonstration. Achievement of the final REE Level One Milestone does not require a flight opportunity. However, the acceptance and infusion of REE technology into other projects will be impaired as a result.

13 Agreements

There are no signed Project agreements as of this writing. A Memorandum Of Understanding (MOU) between REE and the AFRL's ISCP is currently being negotiated. This MOU will cover joint development of software, sensor interfaces, and secondary storage capabilities on the REE 1st Generation Testbed and ISCP architecture. ISCP and REE have independently competitively selected the same contractor for the current phases of each project. The MOU seeks to prevent duplication of work and expansion of the technical development made possible by a common prime contractor.

14 Performance Assurance

The REE Chief Engineer is responsible for performance assurance of all deliverables. The Chief Engineer will employ standard JPL performance assurance processes to test and validate all software and hardware deliverables.

15 Risk Management

Risk can be classified into two general categories: *technical* risk and *resource/schedule* risk. The first refers to uncertainty arising from unexpected development difficulties. The REE Project has been structured to minimize the risk associated with the attainment of Project milestones and their minimum success criteria. While we expect to meet these criteria, there are in addition several "stretch goals," high-payoff/high-risk elements for which success will substantially exceed Project commitments. The second risk category, *resource/schedule* risk, involves factors beyond the control of the REE Project.

15.1 Technical Risk

There are two primary technical risks facing the REE Project:

1) That reductions in power for device component technology will not attain the expected industry projections for the year 2002.
2) That software-implemented fault-tolerance will not prove sufficiently reliable to permit the extensive use of COTS-based technologies.

The impact of (1) could be the failure of the Project to meet the performance criteria for the level one milestones in December 1999 (CT 8) and June 2002 (CT 10). The consequence of (2) is that REE would be forced to include at least some radiation-hardened components in the flight prototype, again lowering performance. In addition, cost would be increased.

The REE Project will mitigate the first risk by making strategic investments in alternative ultra-low power technologies. Several promising, but immature, technologies have the potential for revolutionary breakthroughs in power performance. The key enabling technology for all of these is the dramatic increase in the density of gates that can be implemented on silicon. This trend may permit the placement of fully functional general purpose computers or reconfigurable special purpose computers on a single chip. The elimination of the power normally required to move data off-chip and between chips would represent a significant improvement and could provide REE with an alternative path to the targeted power performance.

The REE Project will mitigate the second risk by leveraging related programs managed by the Air Force and by DARPA. The Air Force Improved Space Computer Program (ISCP) has placed a high premium on system survivability. Consequently, ISCP will invest a major portion of its resources in the development of radiation-hardened components. NASA will coordinate its own milestones and investment strategy with ISCP to leverage this development and provide REE with an alternative path for radiation-tolerance. If necessary, REE will incorporate radiation hardened components in critical sections of the architecture to raise the overall system reliability to the required level. A second mitigation strategy is to incorporate replicated, voted components into the architecture to achieve the required system reliability.

Table 15.1 Technical Risk Assessment.

Risk (probability without mitigation)	**Impact**	**Mitigation**
• Component technologies do not attain power and performance capabilities projected by industry for 2002 (low)	• REE Project fails to meet 300 MOPS/watt success criterion for the flight prototype (Level 1 Milestone CT 10)	• Invest in alternative ultra-low-power technologies, e.g., Processor-in-Memory (PIM) and reconfigurable computing
• SIFT technology does not attain sufficient reliability to permit the extensive use of COTS in space (medium)	• Reliability of flight prototype is compromised. • Performance of flight prototype is compromised • Cost is increased	• Leverage related programs managed by the Air Force and DARPA • Incorporate radiation-hardened components into critical sections of the architecture of the flight prototype • Incorporate replicated/voted components into critical sections of the architecture of the flight prototype

15.2 Programmatic Risk

There are three primary programmatic risks facing the REE Project:

1) That the end result of the REE Project will not be adopted by future NASA missions.
2) That the private sector developers of state-of-the-art software will not allow the REE prime contractor(s) to license and modify their software.
3) That the REE Project could suffer a reduction in the resources available to meet the Project's commitments.

A major concern to the REE Project is that many technology development projects result in technology advances which are not successfully transferred to the intended beneficiaries. This occurs for a variety of reasons, with the primary reason being a lack of attention to the customers needs during the project development. The REE Project is structured to mitigate this risk by engaging the intended customer base (mission science Principle Investigators and mission project managers) from the very beginning of the project. Through the REE Applications Teams, the Project will continuously feed the science missions' needs and requirements into the hardware and software development efforts, so that the end software and hardware technology developed during the Project is driven by and is consistent with the customers' needs for enhanced mission science return at reduced cost.

The consequence of the second risk would be the exclusion of the spaceborne community from the use of popular commercial products, including programming environments, tools, and debugging software. The REE Project will mitigate this risk by minimizing the need to modify COTS software to support software implemented fault tolerance, and maintain active relationships with leading COTS operating systems developers.

Resource reduction is an area of relatively high risk to the REE Project. Annually (and sometimes more often) the Project faces challenges to its budget from all levels of management and oversight. The REE Project has outlined descope options that can accommodate modest resource reductions, while maintaining the overall goals of the Project. In the case of severe reductions, changes to Level 1 milestones will be proposed in a revised project plan. For example, the capturing of a flight opportunity for engineering/science demonstration of REE technology is a "stretch goal." A modest reduction in resources could put this goal at risk. However, the elimination of a flight demonstration does not pose a risk to any Level 1 milestones, since REE does not require a flight in order to satisfy its Level One milestones.

Table 14.2 Programmatic Risk Assessment.

Risk (probability without mitigation)	Impact	Mitigation
• REE technology transfer unsuccessful (low)	• REE technology not adopted by future NASA missions • REE Vision/Goals not realized	• Involve principal REE customer base (instrument scientists) from inception of the project • Continuously feed science-driven requirements into the hardware and software development efforts.

• Private sector developers of software will not allow prime contractor(s) to license or modify their software	• Users of REE technology are precluded from the use of popular commercial software products	• Design SIFT layers to minimize need to modify COTS software • Maintain active relationships with leading COTS operating system developers
• Reduction in funding (high)	• Near-term milestones delayed or descoped with long-term milestones descoped or eliminated • Reduced or elimination of "stretch goals," e.g., flight demonstration.	• Advocate benefits to customers/stakeholders • Re-plan based on project descope priorities

16 Environmental Impact

The Environmental Impact procedures and guidelines are not applicable to the REE Project.

17 Safety

The Safety procedures and guidelines are not applicable to the REE Project.

18 Technology Assessment

The REE Project is a computer research project that pursues technologies that are between five and ten years from maturity. Applications in the areas of Earth and space science are used as drivers of REE's technology research, providing the requirements context for the work that is done.

REE conducts TRL 2–6 research activities intended to prove feasibility, develop and demonstrate computing technologies for eventual introduction into NASA operations though entities such as New Millennium, Discovery, Shuttle and Space Station. REE work in spaceborne COTS parallel computing systems is now at the TRL 2-3 stage, but is planned to attain TRL 6 in 2002.

19 Commercialization

JPL is committed to transferring its technology to the private sector. The following vehicles are available for commercialization of technology, and the REE Project will utilize them depending on mission need and resources.

Technology Affiliates: JPL transfers technology and expertise to U.S. companies on a reimbursable basis to solve key problems identified by the company.

Strategic Technology Development Alliances: JPL develops commercial R&D alliances with U.S. industry focused on shared investment, risk, and benefit strategies.

Targeted Commercialization: JPL targets the commercialization of its validated technologies into emerging global markets.

New Venture Spin-Offs: JPL enables spin-off/start-up companies from the JPL technology base.

Participation in Federal/State Technology Initiatives: JPL establishes a strategic presence in National/State technology initiatives where JPL's technology base will be leveraged for U.S. economic competitiveness and related policy goals.

Regional Economic Growth: JPL encourages economic growth in the region.

In addition, the REE Project will sponsor and conduct technical meetings and workshops and promote the publication of scientific and technical papers to maintain the flow of technology from NASA to industry and academia.

20 Reviews

The REE Project is subject to Independent Annual Reviews (IARs), typically conducted during the last quarter of each fiscal year. These are conducted as part of an overall IAR of the HPCC Program and of the other Projects in the program.

Technical reviews of each Project convened by the HPCC Program are conducted annually. Typically, these consist of end-of-year site reviews at the Project Lead Centers.

The REE Project Manager reports performance monthly to the HPCC Program Office and to the Office of Space Science (Code S).

The REE Project routinely generates the following reports:

REE Project Annual Report

Project Monthly Reports

Each of REE's three element Managers report status on a monthly basis to the Project Manager.

21 Tailoring

The REE Project will be managed and implemented in accordance with the normal procedures used by the Jet Propulsion Laboratory, and in compliance with all requirements established by law and regulations. Executive orders and Agency directives will be observed to the extent accepted by the JPL Prime Contract. There are no major deviations from these procedures.

22 Change Log

May 1997 1st Approved REE Project Plan

Mar 1999 Project Plan revisions to accommodate funding reduction in FY99 and FY00. Plan structure and content revised to conform to NPG 7120.5A and new WBS structure developed to more closely align with the Project major activities

1. Computing Testbeds milestone CT 8 is delayed from 03/99 to 12/99 due to funding reduction in FY99. Low power technology studies suspended for FY99.
2. Grand Challenge Applications milestone GC 6 is delayed from 06/99 to 03/00. The completion of this milestone depends on the completion of CT 8
3. System Software milestone SS 5 is delayed from 03/00 to 09/00. The completion of this milestone depends on the completion of CT 8.

Appendix A
Acronyms

AFRL	Air Force Research Laboratory
API	Application Program Interface
ARC	Ames Research Center
ASIC	Application Specific Integrated Circuit
AT	Aerospace Technology
CAS	Computational Aerosciences
CIC	Computing Information and Communications
COTS	Commercial Off-the-Shelf (Technology)
CPU	Central Processing Unit
CT	Computing Testbeds
DARPA	Defense Advanced Research Projects Agency
DOD	Department of Defense
ESS	Earth and Space Sciences
ESSP	Earth System Science Pathfinder
FLOPS	Floating Point Operations per Second
FPGA	Field Programmable Gate Array
FY	Fiscal Year
GB	Giga (10^9) Byte (of memory)
GC	Grand Challenge
GeV	Giga (10^9) Electron Volts
GFLOPS	Giga (10^9) Floating Point Operations Per Second
GLAST	Gamma-ray Large-Area Space Telescope
GOPS	Giga (10^9) Operations Per Second
GSFC	Goddard Space Flight Center
HPC	High Performance Computing
HPCC	High Performance Computing and Communications
IAR	Independent Annual Review
I/O	Input/Output
ISCP	Improved Space Computer Program (U.S. Air Force)
ITAR	International Traffic in Arms Regulations
JPL	Jet Propulsion Laboratory
KIPS	Kilo (Thousand) Instructions per Second
KOPS	Kilo (Thousand) Operations per Second
MeV	Million (10^6) Electron Volts
MIPS	Million Instructions Per Second
μm	micro (10^{-6}) meter (micron)
MOPS	Millions of Operations Per Second
MOU	Memorandum of Understanding
MPI	Message Passing Interface
NASA	National Aeronautics and Space Administration

NGST	Next Generation Space Telescope
NPG	NASA Procedures and Guidelines
OS	Operating System
OSS	NASA Office of Space Science
PCA	Program Commitment Agreement (Milestone)
PIM	Processor-in-Memory
R&D	Research and Development
RAM	Random Access Memory
REE	Remote Exploration and Experimentation
RFP	Request for Proposal
SEU	Structure and Evolution of the Universe (OSS Theme)
SIFT	Software-Implemented Fault-Tolerance
SS	System Software
TeV	Tera (10^{12}) Electron Volts
TRL	Technology Readiness Level
UPN	Unique Project Number
VNIR	Visible/Near Infrared
WBS	Work Breakdown Structure

Appendix B
REE Project Metrics

This section details metrics that have been established for measuring practical progress toward the REE Project Objectives. These metrics have been developed cooperatively between the Program and Project offices. They will be actively used for evaluation, management, and reporting.

The Project milestones have been constructed to demonstrate steady progress towards the achievement of a practical scalable embedded computing environment for NASA applications. These milestones can be categorized into distinct aspects of the conditions necessary for the Project to be deemed successful: embedded application performance, hardware power performance and usability, system software portability, and overall system reliability. Taken together, achievement of these milestones will constitute de facto achievement of the practical embedded scalable computing environment for space which is the Project's goal. Metrics detailed here will be used to determine when a milestone has been successfully completed and to monitor progress towards achieving each milestone.

The Project milestones express achievements in two broad categories: performance and usability. Each requires different measurement tools and different environment considerations. The performance aspects of milestones are generally straightforward to measure. Usability is more difficult to measure, since the characteristics of usability often are specific to the functionality of a particular piece of software or hardware. Certain general characteristics for high performance systems are necessary conditions of usability and are quantifiable and measurable. These are scalability and speedup. Together with portability, these constitute the primary metrics for the applications and system software.

1	Scalability	reflects the need to execute as large an application configuration as possible in the same elapsed time on different sizes of parallel computing platform configurations with no additional development effort
2	Speedup	reflects the need to execute a specific application configuration in the least amount of time
3	Portability	reflects the practical need for software to be developed on a commercially available system and executed on an embedded system which may not be available until late in the software development cycle or to outlive the effective lifetime of the current generation of HPC systems
4	Power Performance	reflects the stringent limits on power for spaceborne and highly portable or remote earth-based systems. It also reflects the fact that performance requirements may actually be increased from those of non-miniaturized systems.
5	Reliability	reflects the need to have a very high probability of sustained correct operation over long periods of time, including unattended operation in the presence of faults.

These five metrics are defined and the rules for their use are described in the paragraphs below. Their order in the above table does not represent their relative importance, nor will all metrics be applied to every milestone. It is important to note that these metrics have meaning only in terms of a *platform and application* taken together. Hence, these metrics will not be used to rank or evaluate "bare" platforms independent of application software. Instead, evaluations of system configurations will be made using benchmark kernels that are representative of actual REE Project algorithms. Indeed, the primary use of these metrics will not be to evaluate computers at all, but to define success criteria for specific Project milestones, where associated platforms and applications are clearly defined.

❑ **Scalability**

Applications and platforms need to be able to execute efficiently in a variety of configuration sizes without re-engineering. This characteristic is referred to as ***scalability***. This metric is derived from the practical requirement that development costs effectively prohibit either software or hardware from being problem size specific. The economies of high performance computing demand that both software and hardware need to be able to function without change on small, medium, and large problems. Scalability has slightly different meanings when applied to software or hardware.

Software scalability refers to the ability of an application or tool to execute work proportional to platform size with a bounded growth in execution time. This concept is best illustrated by an example. Consider the application of counting the number of zeros in a dataset. Suppose the application takes 10 seconds to accomplish the task for a given dataset size on a single processor. If the application also takes 10 seconds to count the zeros in a dataset 50 times as large on a parallel computer with 50 processors, the software is said to be perfectly scalable. It has no growth in execution time as the problem size is increased proportional to the machine configuration size. As a practical matter, some execution time growth is tolerable if substantially larger applications are enabled. Therefore, the Project will consider software to be scalable if execution time growth for scaled applications is no worse than logarithmic (*ln*) in machine configuration size. Scalability is a dimensionless parameter defined as:

$$S_c(n) = T_{n \times w} / T_w$$

where $T_{n \times w}$ is the execution time for an application which does $n \times w$ work on n processors and T_w is the execution time doing w work on a single processor. Thus, the scalability metric is satisfied if

$$S_c(n) \leq \ln(n)$$

is achieved over a sufficient range of n.

Hardware scalability refers to the ability to assemble functioning platform configurations with the same programming and execution environment and reasonable Mean Time Between Failures (MTBF) in a variety of sizes. The largest configuration size for which these conditions hold is deemed to be the scalability limit. The Project will defer to manufacturers designations of their largest product configuration, absent evidence to the contrary.

❑ **Speedup**

For certain problem classes, absolute time to solution is more important than scalability. ***Speedup*** measures the proportional decrease in execution time for a fixed problem as a function of machine configuration size. Speedup is a dimensionless parameter defined as a ratio of execution times:

$$S_p(n) = t_1/t_n$$

where t_1 is the problem execution time on a single processor and t_n is the execution time on *n* processors. Because a fixed problem, by definition, has a predetermined limit to the number of operations it performs, its speedup will always have an upper bound.

❑ **Portability**

To preserve the value of the initial development investment in an application, ***portability*** of software among the major vendors' platforms is an important attribute of the software design and the execution environment. Portability in the strict sense simply means being able to move an application from one platform to another and have it execute correctly with only a recompile and relink. This implies that the source language(s) is(are) available and that the runtime environment (libraries, OS interfaces, files system interfaces) is the same across platforms. An additional consideration is that the ported application exhibit similar efficiencies (scalability, speedup, performance) on the new platform as on the old. For the purpose of this document, portability is defined as a logical parameter which assume the values "true" and "false." Software which does not require detailed knowledge of the operating system behavior and of hardware configuration will be considered "portable" if it requires no more than name replacements and argument list changes to make it run on a new platform.

For software which requires detailed knowledge of operating system behavior and of hardware configuration, the definition of portability must be relaxed to allow for the construction of custom drivers and interfaces to match the hardware and OS functionality. The software implemented fault tolerance layers will fall into this category. For this class of software, portability will be defined as requiring no more than the replacement of drivers and interfaces totaling less than 10% of the total number of lines of code.

❑ **Power Performance Metrics**

There are stringent limits on the power, mass, and size of systems that are launched into space or developed for highly portable earth-based applications (e.g., lap-top computers). At the same time performance requirements may actually be *increased* over those of past missions (or earlier generation lap-tops). The power performance metrics characterize the ability of a flight system to attain a given performance level per unit electrical power. We will not specifically address the issues of mass, and volume, but expect that commensurate improvements will naturally result from improvement in power performance. In actual experience it is most often the limitation on power that limits performance. Power performance is measured in MOPS/watt, where MOPS is Millions of Operations Per Second (which may be a mixture of 32 bit integer and floating point arithmetic or logical operations). Although MIPS (Millions of Instructions per Second) is a more traditional measure of processor capability, it does not quantify the actual amount of work accomplished on processors which have complex instruction sets. In many cases, though, MOPS and MIPS may be interchangeable.

❑ **Reliability Metric**

Reliability is defined as the probability of "correct operation" up to time t = T given that the system was operating correctly at time t = 0. "Correct operation" is defined as the absence of any fault condition from which the system cannot recover. Partial loss of capability following fault recovery may or may not constitute the loss of correct operation. Reliability can assume values from 0 to 1. Flight subsystem design specifications invariably call for reliabilities very close to 1.

System reliability is exceedingly important for spaceborne applications for the simple reason that a flight computer, once launched, cannot be repaired or replaced. Reliability characterizes the ability of a flight computer to recover from fault conditions (or avoid them altogether), which arise mostly from high levels of radiation. Fault recovery in reliable systems will be achieved with limited loss of performance. Reliability is an overarching metric which encompasses several other familiar attributes of flight computing systems, including fault tolerance and graceful degradation.

❑ **Other Computing Milestone Metrics**

Some milestones require additional metrics which are specific to the milestone. In most cases, these are success counts. They may be specific numbers or a percentage of maximum possible, depending on the milestone, and are indicative of success across a variety of types of applications.

Real time latency is defined as the ability to respond to an external event, such as an externally generated signal to the system, and take appropriate action within a specified period of time. Because the amount of time required to execute a signal handler depends on the details of the handler itself, real time latency is defined here to be the elapsed time between the time an externally generated signal is input to the systems and the time at which the first instruction of the signal handler is executed. This latency is to be accomplished in the presence of faults at rates expected in low Earth orbit.

Each Project milestone will generally require two or more metrics against which progress will be measured. This is due to the complex nature of each of the milestones, and the fact that most milestones require the demonstration of both usability and performance. A milestone will be considered completed when the success criteria for all of the metrics applied have been met or exceeded.

Appendix C
Description of the five REE Applications

Gamma-ray Large Area Space Telescope (GLAST)

Principal Investigators: Prof. Peter Michelson (Stanford University) and Prof. Thompson Burnett (University of Washington)

The Gamma-ray Large Area Space Telescope (GLAST) is a next-generation high-energy gamma-ray telescope that will operate in the energy range from 10 MeV to 300 GeV. GLAST is currently part of NASA's Office of Space Science Structure and Evolution of the Universe (SEU) program strategic plan. The GLAST mission is based on a new pair-conversion telescope design that utilizes modern solid-state particle detector tracking technology (i.e., silicon-strip detectors). To realize the *full* scientific potential of the GLAST instrument will require substantial in-orbit supercomputing resources (about 5 GOPS for the baseline hardware configuration). The two primary areas where supercomputing capabilities can have a major impact on the science return from the GLAST mission are (i) implementation of on-board pattern recognition and event analysis software that will provide the ability to analyze all gamma-ray and cosmic-ray events that trigger the instrument at the hardware level and, (ii) enable real-time analysis of transient events (e.g., the mysterious gamma-ray bursts) and autonomous response to these events. This response could take the form of requests for simultaneous data in real-time from other instruments (earth- or space-based) operating in the x-ray, optical, infrared, or microwave bands. In addition, onboard computing will have a central role in autonomously maintaining instrument calibration and determining alignment of the detector towers.

The availability of supercomputer capabilities in orbit would meet the baseline GLAST computing challenge and would extend the scientific reach of GLAST in important ways. In particular, supercomputing would allow implementation of more sophisticated on-board event triggering and processing that in turn would allow GLAST to (i) measure the energy spectra and elemental abundance of primary cosmic-rays up to some 10s of GeV and measure the flux and energy spectrum of electrons up to the TeV range, (ii) respond quickly to transient events such as high-energy gamma-ray bursts, and (iii) provide the additional computational capability needed to deal with the much larger event size (at least a factor of 5) associated with an imaging calorimeter. The imaging calorimeter can provide additional background rejection capability and enhance the gamma-ray astronomy reach of the instrument above 1 GeV by increasing the effective area at high energies by about a factor of 3 (therefore increasing the rate by a factor of 3 as well). Finally, on-board computing is necessary for monitoring the status of all instrument data channels, maintaining calibration, and determining the relative alignment of the silicon tracker planes. Relative alignment of the tracker channels needs to be known to about 50 μm. This can be accomplished in orbit by using high-energy cosmic-ray proton tracks (that provide straight tracks, relatively free of the effects of scattering) to internally survey the instrument. This alignment calibration will need to be done periodically throughout the mission. Establishing the ability to perform these functions effectively on-board can have important consequences for the actual design and operation of the GLAST. The availability of supercomputer capabilities will enhance instrument performance in all of these areas and has great potential for reducing ground operations cost by reducing the demand for high downlink capacity.

Next Generation Space Telescope (NGST)

Principal Investigator: Dr. John Mather (NASA Goddard Space Flight Center)

In response to the recommendations of the *Hubble Space Telescope and Beyond Committee,* NASA is studying the feasibility of developing a large (8 meter diameter primary mirror) space telescope, optimized for use in the near infrared. The central mission for this instrument, dubbed the Next Generation Space Telescope (NGST), is the study of the early universe: the first stars and galactic structures that are thought to form at redshifts greater than those observable by the Hubble Space Telescope (HST) or other planned facilities. Supercomputer capabilities will have a major effect on the scientific capabilities of the NGST. The two primary areas for investigation are improvements in the data collection from large array detectors with 100 million pixels and improvements in control of the optical system. Improved data collection offers better sensitivity, better immunity to cosmic ray hits, and possibly better calibration accuracy, as well as a reduction in the amount of data to be sent to the ground. Better control of the optical system, which by its nature must be adjusted after launch, could yield better imaging and reduce the overhead of time spent adjusting the figure after it is disturbed. Progress in these areas would have major consequences for the actual design and operations of the NGST. The NGST study has defined a number of stretch technologies which could enable substantial improvements in scientific performance or reduction in cost. Onboard supercomputer capabilities fall in this category.

In the performance of multi-read infrared detector readout and signal processing, large gains in data compression and lowered noise appear possible but will require 100 - 1000 reads per pixel (up to 0.6 Gpixels per sec) and an algorithm to detect and eliminate cosmic rays. The NGST mission is baselined with a primitive version of such a program but larger gains appear possible leading to a reduction in requirements for down link bandwidth and onboard mass storage. With 100 million pixels, even a modest number of samples per second demands a very large compute capability, approaching Gflops or more. The computer memory needs to be large compared with the number of pixels, so at least 1 GB will be needed just for short term fast memory. We do not yet know whether a large memory will be required to hold a long time series for each image, with all 100-1000 reads in memory at once, or whether decisions can be made on the fly so that only a few samples per pixel are kept in the memory.

The availability of an on-board supercomputer will enhance the NGST mission optics in important ways. It will significantly increase the availability of the scientific instruments for scientific observations, by reducing the time required for the periodic fine-figure control. It will improve the quality of the imagery by allowing the adoption of potentially higher-performance closed-loop algorithms for fine-figure control. It would also make possible the adoption of much higher actuator-density deformable mirrors, such as are currently being developed at JPL for coronagraphic imagers. A coronagraphic camera with a second, 10,000 to 20,000 actuator deformable mirror will provide extremely high dynamic range imaging for direct planet detection.

Mars Rover Science

Principal Investigator: Dr. R. Stephen Saunders (NASA Jet Propulsion Laboratory)

NASA has formulated a strategic framework for Mars exploration. The approach is to explore Mars along three thematic lines: search for life, understand climate history, and map resources and geology/geophysics. The strategy is to first obtain global geochemical and mineralogical maps of Mars from orbit. The second step is to characterize and explore sites using rovers that are capable of selecting samples of rock and soil. The third step is to land at one of the previously explored sites, collect a sample and return it to Earth. This strategy will be implemented in a series of missions that include a lander and orbiter in 2001 and in 2003 and the first sample return in 2005. The primary focus is on discovering whether life ever occurred on Mars, and if so, where and for how long. Future robotic missions to Mars, including missions with human crews who will work with robotic field assistants, will use supercomputer capabilities to greatly enhance the scientific return and capabilities of the next generation of Mars mobile platforms.

What is the new science we get with 100 times more computing power? We will develop a plan and partial implementation of software that will make use of 30 - 1000 MOPS/watt, in the range of 150 MOPS to 5 GOPS as compared to the Rover Sojourner at a few watts and perhaps 100 KOPS. The improvements fall into two categories. (1) Navigation: Basically, we want to get from point A to point B faster. The goal is a factor of 10 - 25 faster than Rocky 8 (The 2001 prototype), and access to at least 100 times more area during a mission. (2) Autonomous science operations. (2a) Autonomy to ensure science return in the event of missed commands. The total gain from autonomy is a factor of 6 - 9 in number of fast spectrometer measurements. (2b) Improved science along traverses. The additional science return from opportunistic autonomous observations along a traverse is a factor of about 50 over the return available without REE computing. When compared with the brute-force alternative of launching proportionately more missions to Mars, it is clear that REE computing will be enormously cost-effective for Mars Rover applications.

Orbiting Thermal Imaging Spectrometer (OTIS)

Principal Investigator: Prof. Alan R. Gillespie (University of Washington)

NASA has currently deployed a thermal infrared spectrometer in orbit around Mars to determine surface components for which measurements of reflected sunlight are not diagnostic. Other governmental agencies are actively studying the role that thermal infrared imaging spectroscopy might play in remote sensing here on Earth, and they and NASA are now developing plans for hyperspectral thermal imagers in low Earth orbit. The purpose of these instruments will be not only to collect compositional information, but also to measure land surface temperatures with greater accuracy than has been possible before.

The key impediment to the accurate recovery of land-surface temperature and emissivity data is correction for atmospheric interference with the signal emitted from the land surface. Many approaches have been explored, and the most promising make use of in-scene measurements of the atmosphere rather than external data sources that lack spatial resolution, are taken at different times than the images, and don’t describe the boundary layer just above the surface. It is attractive to estimate atmospheric transmissivity and radiance, and to correct measured radiances for these parameters, at the point of data collection, in orbit.

The Sacagawea satellite proposed to NASA’s Earth System Science Pathfinder (ESSP) was based around a high-resolution HgCdTe imaging system that acquired 64 bands of thermal infrared radiance data at wavelengths from 8.3 to 11.6 μm, with a ground resolution of 30 m, an image swath width of 21 km and a temperature precision of 0.1K. Sacagawea also contained a separate imaging system to measure atmospheric effects at higher spectral resolution, but lower spatial resolution, in the wavelength region 7.5-8.5 μm, and a three-channel Visible/Near-Infrared (VNIR) imager to help distinguish vegetation, clouds and snow. Although Sacagawea itself will not be constructed under the ESSP program, a similar instrument is still under consideration.

On-board processing can be of great benefit in hyperspectral imaging to reduce data volumes and increase duty cycles. The focus of this application will be the development of an on-board processing system to (1) characterize and compensate for atmospheric effects, (2) calculate land surface temperatures and emissivity spectra, and (3) explore automated scene classifiers. In consideration of the diverse user community for these data, transmission of data to Earth may occur at different points in the processing stream. For data that have been completely processed to the thematic map level, data reduction by a factor of ~25 is feasible even before data compression. In extreme cases for which the scene is uniform (large forests, ice caps, water) greater savings are possible.

Atmospheric characterization will make use of a hybrid approach, using a combination of atmospheric data calculated from the atmospheric imager and estimated from forward models driven by climatological and topographic data, all augmented by empirical-line corrections over regions identified as having known surface types and emissivity spectra on the basis of the VNIR data.

Solar Terrestrial Probe Multiplatform Missions

Principal Investigator: Dr. Steven Curtis (NASA Goddard Space Flight Center)

The Solar Terrestrial Probe line of missions was a result of the consensus on the direction of future missions across the Code S enterprise arrived at in Brekenridge, Colorado in 1997. The Solar Terrestrial Probe line is designed to be a series of scientifically linked to pursue a quantitative understanding of the flow of energy, momentum, and mass from the Sun, through interplanetary space, into the magnetosphere, and finally to where it is deposited in the Earth's upper atmosphere. The Solar Terrestrial Probe line is the logical successor to the highly successful International Solar Terrestrial Physics (ISTP) program which has provided the first system level study of the connections between the Sun and the Earth on global scales.

The proposed project will focus on multiplatform missions to study the Sun and the magnetosphere. These missions will consist of by 4 to 100 or more platforms flying in formation. The multiplatform requirement is driven by either image synthesis requirements for remote sensing, for example the low frequency radio imaging of solar processes or the need to uniquely separate space and time for in situ measurements on meso and micro scales, as is the case for the determination of electric currents from the curl of magnetic field variations. Each platform in these missions will have transmission to ground requirements in units of Gbits/day. Since there is an obvious burden on ground systems given the bandwidth requirements, a reduction in the amount of data transmitted to ground is necessary. This can be accomplished by onboard heuristic or high speed data analysis or a combination of both. The focus of this proposal is on the second path.

The tasks chosen for the proposed work are (1) plasma moment calculations for the constellation class nanospacecraft missions presently under study at GSFC as part of the Solar Terrestrial Probe line which are expected to fly in 2007 or later, (2) the calculation of cross correlations between pairs of time series for the imaging low frequency radio astronomy platforms which have been studied at GSFC under the Sun Earth Connections Mission New Concept program and earlier jointly with JPL under the similar Astrophysics New Mission Concepts Program, and (3) the calculation of electrical current from magnetic field variations as measured by a cluster of four or more spacecraft as is being studied both for constellation class missions and for the Magnetospheric Multiscale, the latter of which is expected to fly in 2004 or later.